

Teacher Quality Policy When Supply Matters

Jesse Rothstein*

University of California, Berkeley and NBER

November 2013

Abstract

Recent proposals would strengthen the dependence of teacher pay and retention on demonstrated performance. One intended effect is to attract those who will be effective teachers and repel those who will not. I model the teacher labor market, incorporating ability heterogeneity, dynamic self-selection, noisy performance measurement, and Bayesian learning. Simulations with plausible parameter values indicate that labor market interactions are important to the evaluation of alternative teacher contracts. Reasonable bonus policies create only modest incentives and thus have very small effects on selection. Tenure and firing policies can have larger effects, but must be accompanied by substantial salary increases. Both bonus and tenure policies pass cost benefit tests, though the magnitudes of the benefits are quite sensitive to parameters about which little is known.

1 Introduction

In a 2010 manifesto, sixteen big-city school superintendents stated confidently that “the single most important factor determining whether students succeed in school...is the quality

*Goldman School of Public Policy & Department of Economics, University of California, Berkeley. 2607 Hearst Avenue #7320, Berkeley, CA 94720-7320. E-mail: rothstein@berkeley.edu. I thank Sarena Goodman for excellent research assistance and David Card, Sean Corcoran, David Figlio, Richard Rothstein, Cecilia Rouse, Todd Sorensen, Chris Taber, and numerous conference and seminar participants for helpful discussions. I am grateful to three anonymous referees for their suggestions and to the Institute for Research on Labor and Employment at UC Berkeley for research funding.

of their teacher” (Klein et al., 2010). Influential advocates promise that policies aimed at improving teacher quality can “turn our schools around” (Gates, 2011). As Secretary of Education Arne Duncan puts it, “[w]e have to reward excellence....We also have to make it easier to get rid of teachers when learning isn’t happening” (Hiatt, 2009).¹

A large recent literature focuses on the measurement of teacher effectiveness (e.g., Chetty et al., 2013a; Bill & Melinda Gates Foundation, 2012). Relatively little attention has been paid to the design of policies that will use the new measures to improve educational outcomes. This is an important omission; it is not clear what these policies should look like nor how effective they are likely to be. What evidence exists is discouraging. Several recent experiments have examined the short-term effects of performance bonus policies in U.S. schools, with generally disappointing results (Goodman and Turner, 2013; Fryer, 2013; Springer et al., 2010; though see Fryer et al., 2012).²

Many observers believe that variation in teacher effectiveness primarily reflects personality traits. Klein et al. (2010), for example, urge us to “stop pretending that everyone who goes into the classroom has the ability and temperament” to be an effective teacher. Under this view, the primary mechanism by which instructional quality might be improved is through selection. A well designed contract could make the profession more attractive to effective teachers and less attractive (or unavailable) to ineffective teachers (Lazear, 2003).

This type of effect is difficult to study empirically. Career decisions depend on expected compensation many years in the future, and short-term experimental interventions cannot much affect this. Even quasi-experimental approaches are not promising. Performance pay programs have generally been short-lived (Murnane and Cohen, 1986), so

¹These are not the only potential routes to improved instruction. Alternatives include improved selection at hiring or more or better professional development. Researchers have not identified characteristics observable at the time of hiring that are strongly correlated with subsequent effectiveness (Hanushek and Rivkin, 2006; Clotfelter et al., 2007; Rockoff et al., 2011). However, Taylor and Tyler’s (2012) examination of formative evaluations of experienced teachers found large impacts on teachers’ subsequent performance.

²There is more positive evidence from other countries, mostly poorer than the U.S. See, e.g., Lavy (2002), Muralidharan and Sundararaman (2011), and Duflo et al. (2012); but also see Glewwe et al. (2010).

potential teachers are unlikely to expect that recent pilots will last very long.³

This paper examines the selection effects of alternative teacher contracts. I develop a stylized model of the teacher labor market that incorporates heterogeneity in teaching ability, dynamic learning, and contracts that condition pay or retention on realized performance. Teacher supply responses derive from a dynamic discrete choice model in which graduates and experienced teachers choose between teaching and alternative occupations on the basis of anticipated compensation, which in turn depends on the (potential) teachers' prior information about their own ability.⁴ Decisions to enter teaching depend on risk-adjusted expected compensation over the whole career. Similarly, experienced teachers' exit decisions consider the expected teaching salary over the teacher's remaining career.

A consensus result in the recent teacher quality literature is that characteristics that might be observed before entry into teaching are at best weakly correlated with eventual effectiveness. Thus, I assume that teacher ability is fixed but unknown to either the employer or the teacher herself. Compensation and retention decisions can condition only on a sequence of noisy performance signals, which might be "value added" scores or some alternative. A prospective teacher starts with a private signal about her own ability, then updates her estimate with each performance measure. A teacher who receives positive signals concludes that she is high ability so is likely to receive an above average number of performance bonuses in future years or to have a below average probability of being fired for poor performance, while a teacher who receives negative signals concludes the opposite. These expectations drive the teacher's dynamic decision-making about whether to enter the profession and, having entered, to remain.

³Several large districts, including Denver and Washington DC, have recently adopted performance-dependent contracts. I discuss evidence from Washington in Section 3.

⁴Dynamic occupation choice models include Adda et al. (2013) and Keane and Wolpin (1997); see also Wiswall (2007) on teacher licensing requirements. Essentially static models of teacher attrition include Murnane and Olsen (1989) and Dolton and van der Klaauw (1999). Tincani (2011) develops a static model of teacher sectoral choice under policies like those considered here.

Given the extremely limited variation in extant teacher contracts, I do not attempt to estimate the model. Instead, I simulate the impact of alternative contracts using plausible parameters, exploring the robustness of the results through extensive sensitivity checks.

My policy analysis is closely related to the personnel economics literature on incentive contracts (e.g., Prendergast, 1999; Lazear, 2000) and to studies of teacher firing policies by, e.g., Staiger and Rockoff (2010), Boyd et al. (2011), and Winters and Cowen (2013b,a). The latter studies ignore teachers' behavioral responses. In Staiger and Rockoff's (2010) simulation of tenure policies, for example, teachers can be replaced with new hires without limit, with no consequence for the quality of applicants or the salary that must be paid. Not surprisingly, then, it is optimal to deny tenure to most teachers – 80% or more in the authors' preferred specifications. My model adds a non-degenerate teacher labor market. An increased firing rate requires higher salaries, both to compensate prospective teachers for the increased risk and to attract the needed additional applicants.⁵ I find that the required salary increase is substantial. Optimal firing rates are much lower than those obtained by Staiger and Rockoff (2010). The resulting policies can be cost effective, but offer much net smaller benefits than have sometimes been promised.

My framework allows me to consider a broader class of contracts than do previous studies. I focus initially on a performance bonus based on recent outcomes and a one-time tenure decision, but I also explore contracts that allow for ongoing retention decisions throughout the career. I also examine possible interactions between credentialing requirements – which can be seen as a fixed cost to entering the profession – and contract terms.

In order to focus on the selection margin, I rule out any other effects of performance-sensitive teacher contracts. Effort is irrelevant, all relevant outputs are measured, and teachers cannot influence their actual or measured performance. This neglects the likelihood that

⁵Although layoffs during and after the Great Recession have created considerable slack in the teacher labor market, as recently as 2007 education policymakers worried about where they would find enough qualified new teachers to replace retiring teachers (Chandler, 2007; Gordon et al., 2006).

high-stakes accountability could lead to distortion of the performance measure (Campbell, 1979). Baker (1992, 2002); Holmstrom and Milgrom (1991) find that high-stakes contracts can be counterproductive in this case. I return to this topic in Section 6.

2 What are the Policies of Interest?

In a three-year random assignment study of performance bonuses, Springer et al. (2010) found no effect on student outcomes. But a three year experiment can identify only effects operating through teacher effort. Identifying selection effects experimentally would require that the researcher “start by identifying a couple thousand high school students, follow them for fifteen or twenty years, and study whether alterations to the compensation structure of teaching impacted who entered teaching, how they fared, and how it changed their career trajectory;” even if this could be accomplished, the study “wouldn’t tell us what to do today [and] wouldn’t generate much in the way of findings until the 2020s” (Hess, 2010). Efforts to evaluate selection effects via natural experiments face similar challenges, as past performance contracts have been short-lived (Murnane and Cohen, 1986).

This motivates my structural modeling strategy. A rich enough model can be used to simulate even the long-run effects of policies that have not yet been implemented. The simulation, of course, is conditional on the parameter values used. In principle, the parameters of a correctly specified model could be identified using data on teachers subject to contracts using the traditional “single salary schedule” that ties salaries to education and experience without regard to effectiveness. But results would be highly sensitive to functional form and distributional assumptions. I rely instead on parameter values informed by the best evidence from the literature, and present extensive sensitivity exercises designed to illustrate the sensitivity of the results to the chosen parameter values.

The model has two broad components, a performance measurement system and a

specification of teacher labor supply. A large recent literature examines the former topic, usually through “value added” models that measure a teacher’s effectiveness based on her students’ test score growth. I use estimates from this literature to calibrate the performance measurement component of the model. Nothing in the model is specific to a value-added-based system; it could equally well describe contracts based on more traditional, observation-based performance measures.

There is less guidance in the literature about the parameters governing the labor supply portion of the model. I discuss here several aspects of the policies of interest that bear on this portion of the model.

First, I focus on modeling policy changes implemented at the level of the state or nation rather than by individual districts. This implies smaller labor supply elasticities (Manning, 2005).⁶ Relatedly, I rule out the “dance of the lemons,” a well known phenomenon whereby teachers denied tenure by one district restart their clocks in neighboring districts. In my model, teachers denied tenure must exit the profession.

Second, an important issue in my analysis is uncertainty about a teacher’s ability that is gradually resolved through her demonstrated performance on the job. Accordingly, I model both entry and retention decisions. Because uncertainty is greatest at the beginning of the career, entry decisions are relatively insensitive to the performance component of the contract, while exit decisions become gradually more sensitive as the career goes on.

Third, occupation choices depend on the trajectory of anticipated salaries over the career. While in my model teachers’ expectations incorporate uncertainty about their own abilities and noise in the performance measure, I rule out uncertainty about the future direction of policy (Murnane and Cohen, 1986): Prospective teachers assume the contract

⁶Lazear’s (2000) Safelite Auto Glass study examines a firm-level performance pay program; a similar program implemented at the industry level would likely have smaller selection effects. The ongoing evaluation of the federal Teacher Incentive Fund will assign schools to treatments within participating districts (Glazerman et al., 2011), so at best will identify the partial equilibrium effects of locally-implemented policies.

under consideration will be in effect throughout their careers, and I examine steady-state effects after all teachers recruited under a prior contract have retired.

Finally, teacher quality depends on both supply and demand. I assume districts are unable to distinguish teacher ability at the point of hiring.⁷ Their only options are to adjust contract parameters to induce self selection on the part of potential teachers and/or to retain teachers selectively based on observed performance.

3 The Model

I develop the model in several parts. Section 3.1 defines the performance measure and the Bayesian learning process, Section 3.2 discusses the on-the-job search model that governs entry and exit decisions, and Section 3.3 describes the performance-linked contracts.

3.1 Effectiveness, Performance Measurement, and Learning

Individual i has fixed ability τ_i as a teacher. In the current pool of teachers ability is normally distributed with mean 0 and standard deviation σ_τ , though new contracts may change the selection process and thereby alter that distribution.

A teacher's output depends on her ability and her experience, t , with known return to experience $r(t)$: $y_{it}^* = \tau_i + r(t)$. Each year, a noisy productivity measure is observed by both the teacher and the employer:⁸

$$y_{it} = \tau_i + \varepsilon_{it}. \tag{1}$$

⁷Ballou (1996) hypothesizes that salaries are weakly linked to quality in the cross section because districts facing excess supply do not select the highest-quality applicants. Rockoff et al. (2011) find that information available at the time of hiring is only weakly predictive of effectiveness. My assumption implies that across-the-board salary increases under a single salary contract have no effect on quality (though see Figlio, 2002). This is consistent with a long-run decline in teacher quality as high-ability women's non-teaching options improved (see, e.g., Corcoran et al., 2004), as this reduced relative pay differentially by ability.

⁸In practice, the signal is $\tau_i + r(t) + \varepsilon_{it}$. But both t and the $r(\cdot)$ function are public information.

The noise component, ε , is i.i.d. Gaussian with mean 0 and standard deviation σ_ε . The performance measure is unbiased – all teachers draw their ε s from the same distribution, regardless of who they teach or the methods they use.⁹

Prospective teachers have only limited information about their τ_i s. At entry, teacher i 's prior is $\tau_i \sim \mathbb{N}(\mu_i, \sigma_\tau^2 - \sigma_\mu^2)$, where μ_i represents the teacher's private information and $\mu_i \sim \mathbb{N}(0, \sigma_\mu^2)$ in the population of current teachers. The precision of potential teachers' information can be summarized by $h \equiv V(E[\tau|\mu])/V(\tau) = \sigma_\mu^2/\sigma_\tau^2$, where $h = 1$ corresponds to perfect accuracy and $h = 0$ to a total lack of information. The employer observes neither μ_i or τ_i , so can base compensation and retention only on the sequence of y_{it} s.

Incumbent teachers update their priors rationally as performance signals arrive. The teacher's posterior after t years is

$$\tau | \theta_t \sim \mathbb{N}\left(\frac{t^{-1}\sigma_\varepsilon^2\mu + (1-h)\sigma_\tau^2\bar{y}_t}{t^{-1}\sigma_\varepsilon^2 + (1-h)\sigma_\tau^2}, \frac{1}{t\sigma_\varepsilon^{-2} + (1-h)^{-1}\sigma_\tau^{-2}}\right), \quad (2)$$

where $\theta_t \equiv \{\mu, y_1, \dots, y_t\}$ and $\bar{y}_t \equiv t^{-1}\sum_{s=1}^t y_s$ is the average performance signal to date. I denote the teacher's posterior mean – the first term in (2) – by $\hat{\tau}_{it}$. As t grows, the influence of the original guess shrinks, and $\hat{\tau}_{it}$ converges toward τ_i .

3.2 The Teacher Labor Market: Entry and Persistence

Prospective teachers have von Neumann-Morgenstern utility $u(w)$, defined over annual compensation w , and discount rate δ . A prospective teacher with information μ draws a single non-teaching job offer, providing continuation value ω_1 , from a distribution Ω_1 . She compares this to the utility she will obtain from a teaching career. I denote this by $V_1(\mu; C)$ to emphasize that it may depend both on μ and on the terms of the contract C . She enters

⁹Chetty et al. (2013a) and Kane et al. (2013) argue that real-world performance measures have this property, though see Rothstein (2010) and Rothstein and Mathis (2013).

teaching if $V_1(\mu; C) > \omega_1$.

Each year of teaching represents a single stage of the dynamic decision game. A teacher beginning her t th year, $1 \leq t \leq T$, has information θ_{t-1} . In year t , she receives a performance signal y_t and is paid a salary w_t . This salary may depend on past and/or current performance. The employer then decides whether to offer her continued employment in $t + 1$, again considering her performance to date. $f_t = f_t(y_1, \dots, y_t; C)$ be an indicator for being fired after period t , and let V_{t+1}^f be the continuation value of a teacher who is fired after t . A teacher who is not fired updates her estimate of her own ability based on θ_t and uses this to forecast her future inside earnings and retention probability. She then draws a single outside wage offer, summarized by its continuation value $\omega_{t+1} \in \Omega_{t+1}$, and decides whether to remain in teaching in $t + 1$ or to accept the outside offer.

The value of remaining in teaching in year $t + 1$ is $V_{t+1}(\theta_t; C)$. A teacher whose outside offer ω_{t+1} exceeds this value accepts the offer. Teachers who accept outside offers, either initially or later, can not reenter teaching. V_t is thus defined recursively:

$$V_t(\theta_{t-1}; C) = E \left[u(w_t) + \delta \left\{ f_t V_{t+1}^f + (1 - f_t) \max(\omega_{t+1}, V_{t+1}(\theta_t; C)) \right\} \middle| \theta_{t-1} \right]. \quad (3)$$

The expectation is taken over the teacher's posterior τ distribution following period $t - 1$, as given by (2), and over the distribution of the noise term ε_{it} . Careers end after T periods, so $V_T(\theta_{T-1}; C) = E[u(w_T) | \theta_{T-1}]$.

The Ω_t distribution is calibrated so that the annual exit hazard under the base contract – the probability that ω_t exceeds $V_t(\theta; C_0)$, where C_0 is the baseline, single-salary contract discussed below – is λ_0 and the elasticities of entry and exit with respect to certain, permanent changes in w are η and $-\zeta$, respectively.¹⁰ The Appendix discusses the

¹⁰As $T \rightarrow \infty$ the average career length approaches the inverse of the annual exit hazard, and the elasticity of the career length with respect to the inside wage converges toward ζ . With the parameters used below ($T = 30$, $\lambda_0 = 0.08$, and $\zeta = 3$), the elasticity is roughly $0.77\zeta = 2.3$. The total labor supply elasticity is the sum of the entry and career length elasticities, $\eta + 0.77\zeta$.

distributional assumptions that generate this.

ω_{t+1} is assumed independent of θ_t and τ . The available evidence indicates little relationship between teaching effectiveness and traditional human capital measures (Hanushek and Rivkin, 2006; Rockoff et al., 2011). Several studies find negative correlations between effectiveness and exit from teaching (Krieg, 2006; Goldhaber et al., 2011); given the weak or nonexistent returns to effectiveness in teaching, one would expect the opposite if teaching ability were positively correlated with outside wages.¹¹ Nevertheless, I weaken this assumption later.

An important parameter governing the effect of tenure contracts is V_t^f , the continuation value of a teacher who is fired. I assume $V_{t+1}^f = (1 - \kappa_t) V_{t+1}(\theta_t; C_0)$, where $0 < \kappa_t < 1$ represents the penalty for being fired after t relative to being retained under contract C_0 . (Note that with my assumptions, $V_{t+1}(\theta_t; C_0)$ is invariant to θ_t .) A teacher who exits voluntarily does not pay the penalty. Thus, if κ_t is large, teachers who expect to be fired with high probability will be likely to exit voluntarily beforehand.

Equation (3) does not have a closed-form solution, so I evaluate V_t numerically. The algorithm is described in the Appendix. To simulate the impact of an alternative contract C , I draw teachers from the joint distribution of $\{\mu, \tau\}$, then draw performance measures $\{y_1, \dots, y_T\}$ for each. For each teacher and each year t , I compute $V_t(\theta_{t-1}; C)$ and $V_t(\theta_{t-1}; C_0)$. The ratio of these, along with the labor supply elasticity, determines the probability of entering the profession and, conditional on entering, on surviving from $t - 1$ to t . Note that my assumption of a constant labor supply elasticity allows me to avoid modeling the distribution of teacher types μ in the *population* of potential teachers – changes in the returns to teaching induce proportional changes in the amount of labor supplied by each

¹¹Chingos and West (2012) find that former teachers' salaries are positively correlated with their value-added as teachers. But they also find that value-added is uncorrelated with attrition rates (West and Chingos, 2009). One potential explanation is that τ is positively correlated both with outside salaries and with the individual's taste for teaching, leaving no correlation between τ and the net desirability of a non-teaching offer (at least among those who currently select into the profession). It is the latter that is relevant here.

type that do not depend on the number of people of that type in the population.

3.3 Teacher contracts

The baseline contract C_0 ties pay to experience but not to performance: $w_{it}^0 = w^0 (1 + g(t))$, with $g'(\cdot) \geq 0$. No teachers are fired. Alternative contracts base either w_{it} or retention decisions on the sequence of performance signals to date. I consider two alternatives, performance-based bonuses and performance-based tenure decisions:

Bonus Bonuses are awarded to teachers with high measured performance, averaged across two years to reduce the influence of noise. Thus, in year t bonuses are awarded to all teachers with $\frac{y_{it} + y_{i,t-1}}{2}$ above a threshold y^B ; first-year teachers are ineligible. Total compensation is $w_{it}^B = \alpha^B w_{it}^0 (1 + b * e_{it})$, where e_{it} is an indicator for bonus receipt, b indexes the size of the bonus (as a share of base pay), and α^B is an adjustment to base pay relative to the baseline contract. y^B is set to ensure that in the absence of behavioral responses a share s^B of teachers would receive bonuses each year.

Tenure Teachers are evaluated for tenure after their second years. Any teacher whose average performance to date $\frac{y_{i1} + y_{i2}}{2}$ exceeds a threshold y^F is given security of employment. Those falling short of the threshold are dismissed. y^F is calibrated so that a share s^F of current entrants are tenured. As before, this threshold is fixed; if the ability distribution of new recruits rises, so will the tenure rate. Pay is as under C_0 with an adjustment α^F : $w_{it}^F = \alpha^F w_{it}^0$.

In my model, the optimal pay schedule would have low annual pay and a very large performance-dependent retirement bonus. This is unrealistic, but it is not unreasonable to incorporate information after year 2 into retention decisions. In Section 5.2 I explore the optimal choices of s^B , b , and s^F under the above contract structures, as well as alternative firing rules that better use the available information.

I consider two scenarios for the choices of α^B and α^F . First, I assume that demand is totally inelastic; base wages are set (via α^B and α^F) to yield the same number of teachers as under the baseline single salary schedule.¹² This is consistent with laws and collective bargaining contracts that commonly specify class sizes. Second, I instead assume that the district’s budget is fixed, so that increases in teacher salaries must be offset via reductions in the number of teachers hired. Here, the α parameters are set so that the resulting total labor supply (not counting that which teachers who have been fired would like to supply) will yield a number of classrooms that exhausts the baseline budget at the specified salaries. This is more consistent with the short-run budgeting process in many districts, which treat revenues as exogenous and balance their budgets via so-called “reductions in force.” I assume that balance is achieved over the long run, so that changes in the number of teachers are matched by equivalent savings on facilities and all other expenses.

3.4 Parameter values

My primary simulations set parameters at what I judge to be likely values, attempting to err on the side of optimism about the prospects for performance contracts. These parameters are shown in Column 1 of Table 1. For parameters for which I have less evidence, I also consider a pessimistic scenario (Column 2), in which I expect the performance contracts to be less effective, and an optimistic scenario (Column 3).

I calibrate the productivity measurement parameters using the literature on teacher value-added. The standard deviation of teacher value-added for students’ end-of-year test scores has been widely estimated to be between 0.1 and 0.2, with 0.15 as a reasonable central estimate (e.g., Rivkin et al., 2005; Rothstein, 2010; Chetty et al., 2013a). Many studies have found that teachers improve with experience but level off quickly; I draw $r(t)$

¹²Bonus programs have often been implemented as add-ons to the preexisting base salary. In my framework, this produces excess labor supply.

from Staiger and Rockoff's (2010) estimates for New York City. A number of papers also examine the year-to-year correlation of measured value-added (e.g., Sass, 2008; Bill & Melinda Gates Foundation, 2010). I set σ_ε to produce a reliability ratio for y (defined as $\frac{V(\tau)}{V(y)} = \frac{\sigma_\tau^2}{\sigma_\tau^2 + \sigma_\varepsilon^2}$) of 0.4, near the top of the range identified in Sass' (2008) survey.

The parameter h quantifies the information that prospective teachers have about their own abilities. When Springer et al. (2010) asked experienced teachers, with several years of performance measures under their belts, to forecast their probabilities of winning performance awards, their forecasts were uncorrelated with actual award receipt. This is inconsistent with a large h . Moreover, several studies find that observable teacher characteristics are poor predictors of future effectiveness. The strongest correlations come from Rockoff et al. (2011), who find that a rich vector of academic and personality characteristics explains only 10% of the variance in value-added. This corresponds to $h = 0.1$.

One might expect teachers – who have self-selected into a very secure but low paying occupation – to be unusually averse to risk (Flyer and Rosen, 1997). I use linear utility as a baseline, but also consider in my pessimistic parameters constant relative risk aversion, with coefficient 3, over annual pay.¹³ I use a 3% discount rate in all scenarios.

The outside salary offer distribution is set to yield a constant 8% annual exit hazard under contract C_0 . Careers end after $T = 30$ years. This is roughly consistent with the observed national data, though in these data exit rates are somewhat higher in the first years of teachers' careers; see Appendix Figure 1.¹⁴

Ransom and Sims (2010) examine how teacher retention probabilities vary with the

¹³To my knowledge, mine is the first dynamic occupation choice model to allow for risk aversion. A more complete model would define $u(\cdot)$ over annual consumption and allow agents to borrow and save. Consumption smoothing would make bonus contracts more attractive but would also add considerable complexity. Rothstein (2012) presents an alternative model of risk aversion that captures some income smoothing.

¹⁴An alternative calibration of the outside offer distribution would use salaries earned by exiting teachers. I am not aware of good estimates of the return to teaching experience in non-teaching jobs. If experienced teachers in other occupations are paid like inexperienced workers, policies that raise the turnover rate may dramatically lower the expected lifetime returns to entering teaching. My assumption of a constant exit hazard implies that the return to teaching experience is the same in the outside market as under contract C_0 .

district’s wage schedule. Their estimates imply $\zeta = 1.8$. Ransom and Sims interpret their estimates in a monopsony framework, and many of the teachers in their sample who quit may take jobs in neighboring districts. Their estimate thus likely overstates the occupation-level exit elasticity, though it is not clear by how much.¹⁵ I assume that at the occupation-level $\zeta = 1$, but also consider $\zeta = 0.5$ and $\zeta = 1.5$. Following Manning (2005), I assume the entry elasticity matches (in absolute value) the exit elasticity: $|\eta| = |\zeta|$.

κ represents the permanent effect on future earnings of being fired for poor performance. I assume that a teacher denied tenure in year 2 sees her future earnings reduced by 2%; when I vary the date of firing decisions, in Section 5.2, a teacher fired after t years suffers a $\min\{t, 10\}$ % reduction. This choice of κ is more likely to be too small than too large. By my calculations, Dee and Wyckoff’s (2013) study of performance-based firing threats under the Washington DC IMPACT program implies a firing cost in excess of 85% of base pay¹⁶ In a quite different setting, von Wachter et al. (2009) find that workers displaced by mass layoffs see their earnings decline by 20-30% relative to a control group, with effects that persist for at least 20 years. Laid off workers were often older and displaced from declining occupations and industries. On the other hand, a mass layoff conveys little signal about the worker’s type, where in my model teachers are displaced for cause; moreover, I assume that fired teachers must move to a new occupation and industry.

I assume that base (real) teaching pay rises by 1.5% with each year of experience.¹⁷

¹⁵Clotfelter et al. (2008) study a targeted (but not performance-dependent) bonus program and estimate a school-level exit elasticity between 3 and 4. They do not distinguish exits from the profession, movements to other districts, and movements to other schools in the same district.

¹⁶Under IMPACT, teachers who receive two consecutive “minimally effective” (ME) ratings face dismissal. Dee and Wyckoff (2013), using a regression discontinuity design, find that an initial ME rating increases the annual exit rate by over one-third. As only about 14% of teachers near the ME threshold one year receive an ME the next year (personal communication from Thomas Dee), the implied value of $\zeta\kappa$ is -2.6. The κ in the text is based on $|\zeta| \leq 3$. Dee and Wyckoff’s (2013) parallel analysis of increases in the likelihood of future performance-linked pay increases yields a point estimate of $\zeta = -2.6$, though the confidence interval stretches from below 0 to over 8.

¹⁷Teacher pay is often back-loaded, particularly when pension accumulations are counted. This may serve to lock in mid-career teachers, though empirical exit hazards (see Appendix Figure 1) are non-trivial at all t .

The bonus contract provides a 20% bonus for teachers whose two-year moving average performance exceeds a fixed threshold $y^B = +0.178$, set to ensure that $s^B = 25\%$ of the current teaching workforce would get bonuses. The tenure contract is calibrated to yield a tenure rate of $s^F = 80\%$ given the current ability distribution, corresponding to a threshold of $y^F = -0.167$. Both y^B and y^F are fixed – if the alternative contracts attract more high- τ teachers then more bonuses would be paid or fewer teachers would be fired.

The final parameters are α^B and α^F , the adjustments to base pay under the bonus and tenure contracts. These are calibrated, given the other parameters, either to ensure the same total number of teachers (in steady state) as are obtained under the baseline contract or to satisfy a fixed budget constraint. In the latter calibrations, I assume that a 1% increase in class size would produce savings equal to 3% of the average teacher’s salary but would reduce student achievement by 0.004 standard deviations.¹⁸ Under my baseline parameters, the bonus contract requires a 3.5% reduction in base pay under either demand scenario; the tenure contract requires base salaries to increase by 12.6% to maintain the same number of teachers or by 10.4% to balance a fixed budget. The pessimistic parameters imply higher salaries under each contract and demand scenario, while salaries under the tenure contract are somewhat lower with the optimistic parameters.

¹⁸My cost estimate is based on teacher salaries representing about one-third of total educational expenditures, which I assume are all variable. The assumed effect on achievement is based on Krueger’s (1999) analysis of the STAR class size experiment finds that reducing class sizes from 22 to 15 raised scores by about 0.15 standard deviations; I assume the dose-response function is linear in log class size. Hanushek (1999) and Hoxby (2000) argue that class size has no effect on output. If so, it would be costless to pay for higher salaries via class size increases. I do not consider this possibility here, though in reporting results I distinguish effects coming from class size from those coming via teacher ability and experience.

4 Results

4.1 Noise, information, and incentives

The incentive faced by a teacher i with ability prior $\hat{\tau}_{it}$ depends on the link between this prior and her true ability, the link from true ability to the performance signal, and the link from the performance signal to the contract terms. Moreover, the success of a contract depends on the average incentive perceived by teachers at each true ability level τ_i , among whom there may be much variation in $\hat{\tau}_{it}$. Each of these links serves to dampen the incentives for positive self-selection.

This is easiest to illustrate for the bonus contract. Using the rule of iterated expectations, the average subjective probability as of year t of receiving a bonus in year $t' \geq t + 2$ (so that year- t performance does not directly influence the award) among teachers of true ability τ can be expressed as:

$$E [E [e_{it'} | \hat{\tau}_{it}] | \tau_i = \tau] = E [E [E [E [e_{it'} | y_{it'}, y_{i,t'-1}] | \tau_i] | \hat{\tau}_{it}] | \tau_i = \tau]. \quad (4)$$

The outer conditioning variables can be omitted from inner expectations because the inner variables capture all relevant information – bonus receipt is independent of ability conditional on measured performance, performance depends only on true ability and not on subjective perceptions, and these perceptions depend only on true ability only through $\hat{\tau}_{it}$.

The innermost expectation on the right side of (4) is a step function in $\frac{y_{it'} + y_{i,t'-1}}{2}$, rising from 0 to 1 at y^B . But each of the three outer expectations serves to smooth this out.

First (working our way outward) consider $E [e_{it'} | \tau_i] = E [E [e_{it'} | y_{it'}, y_{i,t'-1}] | \tau_i]$. Using the parameters from Table 1, teachers with τ at the 90th percentile win bonuses only 54% of the time, while those at the 50th percentile do so 9% of the time.¹⁹

¹⁹I express ability in terms of the percentile rank within the *baseline* τ distribution. Of course, this distri-

This is smoothed out further by teachers' uncertainty about their own abilities. At every t , $V[\tau_i | \hat{\tau}_{it}] > 0$. Because $E[e_{it'} | \tau_i]$ is convex, this uncertainty flattens the $E[e_{it'} | \hat{\tau}_{it}] = E[E[e_{it'} | \tau_i] | \hat{\tau}_{it}]$ function – even teachers who think they are likely to be of low ability realize that they might in fact be of higher ability and thus be eligible for bonuses, while those who think they are of high ability harbor doubts about this so are not confident that they will earn bonuses. This is particularly true of early career teachers, for whom $V[\tau_i | \hat{\tau}_{it}]$ is large. Even a prospective teacher at the 90th percentile of the μ distribution thinks she has only a 37% chance of receiving a bonus in any given year of her career, while a prospective teacher with 10th percentile μ anticipates a 4% chance. As teachers accumulate information, they quickly learn their places in the distribution. After one year, the teacher at the 90th percentile of the $\hat{\tau}_{i1}$ distribution thinks her chance of receiving a bonus is 42%, and this rises to 45% after two years and 48% after 5 years.

The $E[e_{it'} | \hat{\tau}_{it}]$ curve governs the incentives faced by teachers of different $\hat{\tau}_{it}$ s. But it does no good to attract teachers who mistakenly believe themselves to be of high ability and repel those who underestimate their own τ s. The degree to which the bonus contract attracts teachers who are actually of high ability depends on $E[E[e_{it'} | \hat{\tau}_{it}] | \tau_i]$. Because the slope of $E[\hat{\tau}_{it} | \tau_i]$ is less than one, this further attenuates the incentives. Again, this is particularly true early in the career: At entry, the average 90th percentile teacher's subjective expectation of her own ability puts her at only the 70th percentile.

The solid line in the left panel of Figure 1 shows the average probability of winning a bonus by *true* ability percentile, $E[e_{it'} | \tau_i = \tau]$. The other series in this panel show average *subjective* expectations among teachers of each true ability, $E[E[e_{it'} | \hat{\tau}_{it}] | \tau_i = \tau]$, at several points in the career. At the point of entry into teaching there is relatively little differentiation except at the extreme tails of the distribution: The average 90th percentile teacher anticipates a 30% of earning a bonus in any given year while the average 10th percentile

bution would change under alternative contracts. The fixed-norm percentiles are simply a convenient scale.

teacher perceives a 9% chance. But perceived incentives become much better targeted as teachers gain experience – by the 5th year of their careers, the 90th percentile teachers perceive their chances at 44%, on average, while the 10th percentile teachers see theirs as under 2%. Thus, while incentive effects of a bonus system are weak at the recruitment stage, later attrition decisions may be more sensitive to these incentives.

The right panel of Figure 1 repeats the exercise for the tenure contract. (I omit the curve for 5th year teachers, as tenure decisions have been made by then.) Again, we see weak incentives for low ability potential teachers to select other careers at the entry point, but after even a single year the incentives are stronger.

There is a close, albeit imperfect, mapping from the subjective probabilities of positive and negative outcomes graphed in Figure 1 to the average values of teachers of different abilities under the two contracts. Figure 2 shows average continuation values V of teachers under the two contracts, by ability level and years of experience.²⁰ Because the V scale is not intuitive, I convert it to equivalent variations: Changes in w^0 that would yield the same values under the single salary contract. An estimate of +5% means that w^0 would need to rise by 5% under contract C_0 to yield the same value as is obtained under the alternative contract. The figure shows that the bonus contract produces the equivalent of a 0.3% salary increase for the average 90th percentile teacher at entry, and a 4.9% increase after five years.²¹ These changes are not large, and suggest that any self-selection responses to the bonus contract will be quite modest. The tenure contract achieves a much steeper slope among first year teachers, with a range of about 10% of baseline salaries between 10th and 90th percentile teachers, but like the bonus contract does little to alter the rela-

²⁰Averages are computed over all teachers who enter under the baseline contract, ignoring voluntary exits. I assume that base salaries have been adjusted to preserve the size of the workforce. Note that the probabilities in Figure 1 depend only on the performance measurement parameters, while those in Figure 2 depend also on the labor supply and outside wage offer parameters.

²¹Recall that base salaries are adjusted by a factor α^B under this contract. The equivalent variation thus approaches $\alpha^B - 1 = -3.5\%$ as the (subjective) probability of future bonus receipt approaches zero.

tive incentives governing entry to teaching. After the tenure decision teachers benefit from the increased salaries needed to attract sufficient applicants, and values reflect the 12.6% increase in salaries across all ability levels.

4.2 Impact of incentives

I next turn to the impacts of the contracts on the teacher ability distribution. Figure 3 shows how the two contracts would change the number of entering teachers at each ability percentile, relative to the baseline contract. The bonus contract entices more high ability and fewer low ability entrants, but the impact is extremely small. The tenure contract, with its sharply increased salaries, attracts more of all types, but again does little to change the relative numbers of high and low ability teachers.

Figure 4 shows the effects of the two contracts on average career length. The bonus contract has small effects on this margin, too, concentrated at the very top of the ability distribution. The $1 - \alpha^B = 3.5\%$ cut to baseline salaries under this contract reduces career lengths of below-median teachers by only about 2%.

The tenure contract has a more dramatic effect. Career lengths shorten by as much as 80% for the weakest teachers. This primarily reflects tenure denials. Among slightly higher ability teachers, around the 20th percentile, career lengths shorten by an average of about 25%. This change would be less than one third as large in the absence of labor supply responses; the bulk of it derives from voluntary exits after the first year among teachers who have learned their tenure chances are lower than anticipated but who would nevertheless receive tenure if they stayed. At the other end of the distribution, careers lengthen by nearly 10%, as higher salaries reduce voluntary attrition among tenured teachers.

Figure 5 presents the impact of the two contracts on the steady state number of teachers at each ability level, combining entry and career length effects. Not surprisingly,

the bonus contract has little effect, reducing the number of low ability teachers by about 3% and increasing the number of high ability teachers by a bit more than this. The tenure policy is much more effective, increasing the number of classes taught by top-quartile teachers by about 20% and reducing those taught by bottom decile teachers by 60%.

Columns 2 and 3 of Table 2 show the effects of the two contracts on teacher ability, experience, effectiveness (combining ability and experience effects), and salaries. The bonus contract yields only small increases in average teacher ability, around 2% of a teacher-level standard deviation, while the tenure policy has over eight times the effect. The tenure contract increases the number of first year teachers by about an eighth, but this does not have much impact on overall productivity – first year teachers are about half a standard deviation less productive than experienced teachers, so a 0.9 percentage point increase in inexperienced teachers reduces average productivity by less than 1/200th of a standard deviation. Thus, net effects on teacher effectiveness – +0.004 student-level standard deviations for bonuses and +0.033 for tenure – are the same to three digits as the gross effects on teacher ability. Total salary costs, inclusive of bonuses, are essentially unchanged under the bonus contract but rise by 15% under the tenure contract.

Columns 4 and 5 present results for the alternative scenario in which the total district budget, rather than the number of teachers, is fixed. Here, increased per-teacher costs are offset by reducing the number of teachers per student. The tenure contract requires increasing class size by 3.4% to finance the higher salaries that it requires. (These are lower than in Column 3, as the district's labor demand curve shifts inward and it thus can move down the supply curve.) This reduces student achievement, offsetting a bit less than half of the gains from improved teacher quality. The net effect of the policy is to raise average output by 0.018 student-level standard deviations, or about 12% of the baseline standard deviation of teacher productivity. That this is positive indicates that the tenure contract is cost-effective relative to class size reduction under the parameters used for these

simulations. The bonus contract is as well, but its impact is less than one-quarter as large.

4.3 Sensitivity to alternative parameters

The results in Table 2 reflect the parameter values from Column 1 of Table 1. Table 3 presents estimates for alternative parameter values. Here and hereafter, I use the “fixed budget” demand scenario, as this allows me to summarize the impact of alternative contracts by the constant-budget impact on average output, including class size effects. The entries in the first row of Columns 1 and 4 repeat the estimates from Table 2. Subsequent rows vary the different parameters in turn, one at a time. To illustrate potential interactions among parameters, Columns 2 and 5 show results for the “pessimistic” parameter values from Table 1, while Columns 3 and 6 show results for the “optimistic” parameters. (Blank cells correspond to parameter values that match those used for Row 1 of the same column.)

Working down Columns 1 and 4 we can see the impact of each parameter on the results. The impacts of the two policies are not very sensitive to changes in the reliability of prospective teachers’ private information (Rows 2-3), but would grow noticeably if the performance measure could be made more reliable (Row 4).

Rows 5 through 8 show the effects of varying the labor supply elasticities. Both policies are more effective when labor supply is more elastic, particularly on the exit margin. The relative unimportance of the hiring elasticity reflects the fact that entering teachers have too little information about their own abilities to perceive large changes in incentives.

Row 9 shows a variant in which the exit elasticity is allowed to decline with experience, starting at 1.5 and falling to 0.5 by year 10. This is meant to capture the intuition that early career teachers may be more mobile than are later career teachers (Ransom and Sims, 2010). This has nearly as large an effect as does raising the elasticity throughout the career, suggesting that it is early career attrition that most affects the impacts of the policies.

Row 10 shows an additional variant in which the entry elasticity is increasing with μ : $\eta = 1.01 + 0.205 \frac{\mu}{\sigma_\tau}$. In a Roy model of career choice with $\text{corr}(\mu, \omega_1) > 0$, entry of higher- μ potential teachers is more sensitive to the offered wage than that of those with lower μ . The function here is approximately what would obtain with a correlation of 0.1 and a log-normal outside wage distribution calibrated to yield an average elasticity of 1. Allowing for this sort of heterogeneity in entry elasticities has little impact on the results.

In Row 11 I assume that agents are risk averse, with constant relative risk aversion parameter 3, over annual incomes. This has only a minor effect. Row 12 shows estimates for larger firing costs ($\kappa = 0.15$). This reduces the benefit of the tenure policy somewhat under the baseline parameter vector, but has a much smaller effect under optimistic parameters. Finally, Row 13 uses a lower baseline exit hazard, especially for experienced teachers. Appendix Figure 1 shows that many teachers exit the classroom to work in school administration (e.g., as principals). It is not clear that these should be treated as exits for better offers (nor is it clear that they should not be). This change does not much affect the results.

Changes in several parameters at once can be examined by comparing across columns. Column 5 shows that the +0.018 net benefit of the tenure policy turns slightly *negative* under the pessimistic parameters. This largely reflects the reduced exit elasticity ($\zeta = 0.5$, vs $\zeta = 1$ in the baseline). When I adjust the exit elasticity to $\zeta = 1.5$, keeping all other parameters as in the pessimistic scenario (Row 8), results are quite similar to those seen in the baseline scenario with $\zeta = 1$ but more favorable values for all other parameters.

5 Alternative policies

In this Section, I broaden the policy space beyond the simple bonus and tenure contracts considered above. First, I consider combining these policies with reforms designed to make it easier to dip one's toe in the profession. Second, I present estimates for quantitative

changes to the policies, aimed at identifying how the optimal design of each policy varies with the model parameters. Third, I consider alternative retention policies that make better use of the available information than does a once-and-for-all tenure review.

5.1 Interactions with credential requirements

Employment as a teacher traditionally requires earning a teaching credential. This is controversial: It is not clear that credentialing programs provide useful training (Boyd et al., 2006; Kane et al., 2008), and the requirement may serve as a barrier preventing some potentially able teachers from entering the profession. Credential requirements may interact with the policies considered here by making it more difficult to try out teaching for a short time before pursuing other careers, impacting the tenure contract in particular.

To explore this interaction, I augment the model by requiring prospective teachers to pay a fixed cost, equal to a year's salary under the baseline contract, before entering the profession.²² I then consider eliminating this cost, either alone or in combination with the adoption of a performance-based contract. I assume that teacher candidates learn nothing in credentialing programs and that these programs serve no filtering function, so the requirement offers no benefits and its elimination comes as pure gain.

Results are presented in Table 4. Column 1 repeats results for the baseline case of no entry costs considered above. In Column 2, the first rows show the effects of introducing performance contracts when there is a fixed cost to entry. These are identical to the baseline results. The third row of the table show the effect of eliminating the fixed cost, without other changes to the contract. This allows salaries to be lowered, freeing up enough money to finance 1.7% reductions in class size and thereby to improve productivity by 0.006 student-

²²There are aspects of teaching other than credential requirements – e.g., the need to invest at the beginning of the career in the development of lesson plans that may be reused later – that can also be interpreted as fixed costs to entering the profession. The discussion here applies to these costs as well, with an important distinction: Unlike with certification requirements, it is unclear how these costs could be reduced.

level SDs. This is about 50% larger than the impact of the bonus contract.

Finally, the last rows show the effects of simultaneously eliminating the entry cost and introducing the performance contracts. These are equal to three digits to the sum of the separate effects of the two components in isolation. This is not what might have been expected. When compensation is backloaded, as it is with a fixed entry cost and stable growth in post-entry earnings, tenure denials are more costly, and one might expect that teachers – particularly low ability teachers – would demand larger compensation for accepting this risk. Intuition comes from the information structure illustrated in Figure 1. Prospective teachers do not have enough information to forecast their tenure probabilities accurately, so the prospect of a tenure denial weighs nearly equally on high- and low-ability prospective teachers. Both elimination of the entry cost and introduction of the tenure policy thus have effects on entry that are largely uniform across the ability distribution.

There are two important caveats. First, I assume risk neutrality. The policies might interact meaningfully if teachers were risk averse.²³ Second, calls for credential reform are often motivated by the idea that many high-skilled graduates who foresee highly-paid professional careers could be persuaded to teach for a short time if entry costs were low.²⁴ In my model, these graduates have no higher τ than do those with worse outside options, so there is little benefit of attracting them. Allowing for a correlation between inside ability and outside options would make credentialing reform more attractive, and might also cre-

²³Above, I incorporated risk aversion over annual income. This is unsatisfactory in the presence of a fixed entry cost, which prospective teachers presumably finance out of later earnings. Combining risk aversion with student debt requires a more elaborate model of intertemporal decision making.

²⁴Teach for America (TFA) is an example: TFA teachers are asked to commit to only two years of teaching, and are assigned classrooms after a 5-week training course. Clark et al. (2013) find that TFA teachers are more productive than their more experienced peers. This could indicate a pool of high- τ potential teachers who are unwilling to pay fixed costs to enter. It could also reflect, however, a TFA screening process that successfully selects on τ : Only 20% of applicants are selected. In any case, it is not clear that the pool of high-ability potential short-term teachers is large enough to scale up dramatically. After years of rapid growth, TFA now produces less than 0.3% of all public school teachers. And “no excuses” charter schools (Abdulkadiroğlu et al., 2011) that hire from similar pools serve only a few percent of urban students but are already importantly constrained by labor supply shortages (Wilson, 2009).

ate an interaction between credential and contract reforms: Short-term teachers would not be much affected by tenure decisions, so the salary increases needed to offset the risk of tenure denial for a career teacher would make the profession more attractive to short-term potential teachers. I defer consideration of this to future work.

5.2 Varying the bonus and tenure contracts

The bonus and tenure policies I have considered thus far are designed to resemble in both their forms and scopes policies that have been implemented by states and school districts. They are far from optimal. In the context of the model here, the optimal pay policy would backload nearly all compensation to the teacher's retirement, when the teacher's ability is known with maximal precision. This is unrealistic for reasons beyond the scope of my stylized model: Teachers have consumption needs that make it impossible to wait many years for their salaries, and even if credit were available the risk that one will turn out to have low τ and thus never be paid would not be easily insurable.

It is thus not useful to investigate the first-best policies, and in any case computational costs would be prohibitive. I examine here alternatives to the bonus and tenure policies that may be closer to optimal, while maintaining the same basic structure; in the next subsection, I examine a plausible changes in the structure of retention decisions. I assume throughout that the district's budget is fixed and that changes in salaries per teacher are offset by changes in the number of classrooms, as in Columns 4 and 5 of Table 2.

The left panel of Figure 6 shows how the impacts of the bonus contract vary with the size of the annual bonus (expressed as a share of base pay). Not surprisingly, when teachers are risk neutral the effectiveness of the bonus policy increases with the size of the bonus. But the impact of the bonus policy remains small: Under baseline parameters, even a bonus equal to 100% of base salaries would have a smaller impact than would a policy of

denying tenure to 20% of second-year teachers. When teachers are risk averse, as with the pessimistic parameters, large bonuses are counterproductive. Here, the maximum impact is achieved with a bonus equal to 21% of base pay.

The right panel of Figure 6 varies the threshold for receiving the bonus (expressed as f^B , the share of current teachers who would receive bonuses each year). Regardless of parameter values, the impact of the bonus policy grows with the share of teachers receiving bonuses until this share exceeds 40%.

Figure 7 turns to the tenure contract. Here, I vary both the share of teachers denied tenure and the date at which the decision is made; different panels correspond to the different parameter vectors. The upper left panel shows that optimal tenure denial rates under the baseline parameters are around 40%. The tenure policy is notably less effective if tenure decisions are made after only one year, but there is little benefit (or cost) of waiting more than two years. The lower panels show results for the pessimistic and optimistic parameters. Under the pessimistic parameters (lower left), a 40% denial rate is far too high; the optimum is less than 15%, and high tenure denial rates are worse than the status quo. Kater decisions are preferable here. Under optimistic parameters, by contrast, the optimal policy denies tenure to well over half of new teachers, and benefits are not very sensitive to the date of the decision so long as it is made in the second year or later.

The optimal bonus and tenure policies are characterized in the first and second panels of Table 5. Tenure policies are uniformly more effective than bonus policies of plausible scale, but the design and impacts of these policies depend critically on the parameters. Even the optimistic parameter values suggest that the benefits of a tenure contract top out around 0.038 student-level standard deviations. This is less than half of the impact suggested by Staiger and Rockoff's (2010) simulation of optimal tenure policies without labor supply responses. Moreover, where Staiger and Rockoff (2010) estimate that tenure decisions should be made after just one year and over 80% of new teachers should be dismissed, my

results point to later decisions and much higher tenure rates.

5.3 Alternative firing policies

Once-and-for-all retention decisions make inefficient use of information: For teachers whose initial performance places them near the retention threshold, it would be better to extend the probationary period. It is computationally infeasible to solve for the first-best optimal retention rule taking labor supply responses into account. Instead, I consider here three alternative firing contracts that successively better approximate the optimal decision rule in the absence of labor supply responses. The contracts vary in the way that the retention threshold varies with teacher experience.

My first ongoing firing contract conditions retention on the teacher's average performance to date: Any t th-year teacher for whom $\bar{y}_t \equiv \frac{1}{t} \sum_{s=1}^t y_s$ falls below a fixed threshold is fired. I express the threshold in terms of the share of entering teachers who would be displaced at some point before the end of a 30-year career (assuming that they stay that long, and given the baseline ability distribution. For example, a threshold of -0.26 would lead 20% to be displaced at some point before the end of a 30-year career. If the contract were implemented suddenly, 6.8% of teachers would be fired immediately, but many of these are teachers who would have been displaced much earlier had the contract been in place before. Moving forward, 14% of new hires would be fired after their first years; 3.1% of those who remain would be fired after their second years, and 1.3% after their third years. Firing rates would be below 1%, but always positive, for $t > 3$.

This contract is in some ways more patient than the tenure contract – 17% of firing decisions take place after the second year. But it is not patient enough. There is more value in retaining a first year teacher with $y_1 = -0.3$ than in retaining a 10th year teacher whose average performance to date is so poor, as there is a reasonable prospect that the

former teacher was simply unlucky but the latter's performance more likely reflects her true ability. My second ongoing firing contract bases retention decisions on the district's posterior mean of the teacher's ability, $\frac{\sigma_{\tau}^2}{t^{-1}\sigma_{\epsilon}^2 + \sigma_{\tau}^2} \bar{y}_t$. When this contract's threshold is set at a level that would lead 20% of teachers to be displaced at some point in their careers, 3.7% of teachers (given the current ability distribution) are fired after their first years, 4.4% of those who remain are fired after their second years, and 3.1% after their third years. Firing probabilities remain above 1% through the 6th year.

Even this might be too impatient. The option value of retaining an inexperienced teacher with low posterior mean but high variance is higher than for an experienced teacher with the same posterior mean (so better average performance to date) but low variance – the inexperienced teacher may turn out to be fine, and can always be fired next year if she doesn't. My third contract uses thresholds that vary over time in a way that is optimal from the district's perspective, ignoring labor supply responses.²⁵ This contract displaces only 1.3% of first year teachers, 2.5% after their second years, and 2.1% after their third years.

Appendix Figure A.2 shows the share of teachers at each ability level who are eventually fired under each of the different contracts, when each is calibrated to displace 20% of teachers at some point over a 30 year career. The first firing contract does slightly better than the tenure policy at identifying the teachers with the lowest true ability for firing, but the difference is relatively small. The second and third contracts represent more dramatic improvements, firing many more of the bottom quintile teachers and many fewer teachers outside the bottom quintile. However, this comes at a cost. Appendix Figure A.3 shows the cumulative firing probability for teachers in the bottom decile of the τ distribution, the second decile, the third and fourth deciles, and the upper six deciles. Although the more

²⁵The thresholds are computed as the solution to the district's dynamic optimization problem, assuming that the district pays a firing cost that is proportional to the number of years of labor supply foregone and ignores labor supply responses. Note that I fix the overall firing rate by setting the firing cost; if it is set above the district's true shadow cost, the firing rate is suboptimally low.

patient contracts eventually fire larger shares of the lowest ability teachers, they wait longer to do so – substantially so for 2nd decide teachers under the “optimal” decision rule.

Figure 8 presents the impacts of the alternative decision rules at different scales. It shows that the ongoing firing contracts support higher firing rates but achieve only slightly larger impacts than does the tenure policy. As expected, the most patient contract has the largest impacts when the firing rate is set to maximize output, but when the firing rate is kept below its optimal level it can be better to make faster, more error prone decisions than to wait to optimally distinguish among teachers just above and below the desired threshold.

Results for the optimal scale of each contract type are presented in the lower panel of Table 5. All but one of the ongoing retention contracts outperforms the tenure contract; in each case, more patience allows larger net impacts than do the less patient contracts, usually but not always with higher firing rates. However, the optimal decision rules and policy impacts are quite sensitive to the model parameters. Where under the pessimistic parameters the firing rate never exceeds 18%, the baseline parameters yield optimal firing rates as high as 55%, and the optimistic parameters yield firing thresholds that would displace as many as 71% of teachers before the end of a 30-year career. The alternative contracts would yield net productivity improvements ranging from just over 2% of a teacher-level standard deviation, for the average performance contract under pessimistic parameters, to nearly 38%, for the most patient decision rule under the optimistic parameters.

This is much smaller than the effects, well above one standard deviation, assumed in Chetty, Friedman, and Rockoff’s (2013b) analysis of a policy of firing the bottom 5% of teachers based on estimated ability. Another point of comparison is Staiger and Rockoff’s (2010) more thorough analysis of teacher firing decisions. One of the contracts considered there would allow a teacher to be fired anytime over the first four years, under a similar information structure to that used here but without labor supply responses. Staiger and Rockoff find that decisions should be made very quickly, with the vast majority of firings

occurring after the first year; that over 80% of teachers should be dismissed; and that the optimal contract yields net benefits roughly twice as large as those obtained with even my most optimistic parameter values. This points to the important constraints that labor market responses impose on the use of performance-based contracts to improve teacher quality.

6 Discussion

The simulations presented here suggest that the effects of policies aimed at improving teacher productivity will depend importantly on their interaction with the teacher labor market. If prospective teachers are uncertain about their own abilities or if their labor supply is less than perfectly elastic, both performance-based compensation and retention policies require substantial increases in teacher compensation.

Despite the high costs, both classes of policies are cost effective at modest scales in my model. Indeed, recognition of the labor market effects can make non-retention policies even more effective than when these effects are ignored, as the accompanying salary increases help to attract and retain high ability teachers. My results also point to the importance of policy design, as cost-effectiveness varies with the specifics of the teacher contract.

There are several important caveats, however. First, the results of course depend importantly on my choice of parameter values. Policies that are optimal under one set of parameter values can be harmful under other plausible parameters. Results are particularly sensitive to the labor supply elasticity and the degree of foreknowledge that prospective teachers possess, and future research should aim to uncover these parameters.

Second, the analysis relies on a best case view of the potential for teacher performance assessment. I assume that performance measures are unbiased, cover the full range of desired outputs, and are not subject to “influence activities” that raise measured performance without raising true productivity. None of these is very plausible.

Consider first the case where output is multi-dimensional, and the performance measure captures only one of the dimensions. For example, the performance measure might focus on cognitive skills where teachers also teach non-cognitive skills, or it might focus only on certain subjects, or weight test-taking skills too heavily relative to others. Then the policies I consider here will improve teacher ability on each dimension in proportion to the correlation in ability across dimensions. What evidence there is suggests that this correlation is low. For example, several studies find that teacher value-added measures based on one student test are not highly correlated with measures constructed from a different assessment (even when both test achievement in the same subject). Typical correlations are in the 0.4 range, after adjusting for attenuation due to sampling error (see, e.g., Bill & Melinda Gates Foundation, 2010; Rothstein, 2011). Correlations between value-added scores and other forms of performance measurement (e.g., classroom observations) are even lower (Bill & Melinda Gates Foundation, 2012; Rothstein and Mathis, 2013).

Matters are even worse if teachers have some ability to affect their measured performance, either via reallocating effort between measured and unmeasured outputs (also known as “goal distortion”) or by manipulating the performance measurement process. High-stakes evaluations can be counterproductive in this case (Baker, 1992, 2002; Holmstrom and Milgrom, 1991). There is evidence that teachers can improve their measured performance under value-added-based evaluations by reducing the attention paid to non-tested topics and subjects, by teaching to the test, by arranging to have the right students, or by outright cheating, and that teachers faced with high-stakes incentives will respond at least in part in these ways (e.g., Campbell, 1979; Neal and Schanzenbach, 2010; Rothstein, 2010; Carrell and West, 2010). Rothstein (2012) finds that the benefits of performance-based contracts are quite sensitive to the potential for goal distortion and manipulation. A high priority topic for future research must be the degree to which productivity measures become corrupted when the stakes are raised (Rothstein, 2011).

Finally, there are many aspects of the teaching profession omitted from my stylized model. I do not account for the possibility that teachers may be self-selected for unusual risk aversion; for the social status of teachers relative to other professions; or for the potential for high-stakes evaluations to undermine cooperation among teachers and principals. Moreover, I assume that new teachers recruited under alternative contracts would come from the same general population as do current teachers and do not allow for the possibility, sometimes raised in discussions of teacher quality, that there exists a separate pool of high ability potential teachers who would not consider teaching under current conditions. These issues are not well enough understood to incorporate into my quantitative model; extensions of the model to allow for them are left as a subject for future research.

These caveats aside, the analysis here demonstrates that clear thinking about the impact of teacher quality policy changes requires a model of the roles of imperfect information and teacher labor supply decisions. Even in my best case scenarios, alternative teacher contracts have more modest impacts on student achievement than has often been promised. None of the contracts considered here would raise average productivity by more than 40% of a standard deviation. More plausible parameters and policies yield improvements that are generally less than half that size. These kinds of benefits would be most welcome, but would not represent fundamental changes in our education system.

References

- Abdulkadiroğlu, Atila, Joshua D Angrist, Susan M Dynarski, Thomas J Kane, and Parag A Pathak**, “Accountability and flexibility in public schools: Evidence from Boston’s charters and pilots,” *The Quarterly Journal of Economics*, 2011, 126 (2), 699–748.
- Adda, Jerome, Christian Dustmann, Costas Meghir, and Jean-Marc Robin**, “Career Progression, Economic Downturns, and Skills,” Working Paper 18832, National Bureau of Economic Research, February 2013.

- Baker, George P.**, “Incentive Contracts and Performance Measurement,” *Journal of Political Economy*, 1992, 100 (3), 598–614.
- , “Distortion and Risk in Optimal Incentive Contracts,” *The Journal of Human Resources*, 2002, 37 (4), 728.
- Ballou, Dale**, “Do Public Schools Hire the Best Applicants?,” *The Quarterly Journal of Economics*, 1996, 111 (1), pp. 97–133.
- Bill & Melinda Gates Foundation**, “Learning About Teaching: Initial Findings from the Measures of Effective Teaching Project,” Research Report, Bill & Melinda Gates Foundation, Seattle, Washington 2010.
- , “Gathering Feedback for Teaching: Combining High-Quality Observations with Student Surveys and Achievement Gains,” Research Paper, Bill & Melinda Gates Foundation, Seattle, Washington 2012.
- Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff**, “Teacher Layoffs: An Empirical Illustration of Seniority Versus Measures of Effectiveness,” *Education Finance and Policy*, 2011, 6 (3), 439–454.
- , **Pamela Grossman, Hamilton Lankford, Susanna Loeb, and James Wyckoff**, “How Changes in Entry Requirements Alter the Teacher Workforce and Affect Student Achievement,” *Education Finance and Policy*, 2006, 1 (2), 176–216.
- Campbell, Donald T.**, “Assessing the Impact of Planned Social Change,” *Evaluation and Program Planning*, 1979, 2, 67–90.
- Carrell, Scott E. and James E. West**, “Does Professor Quality Matter? Evidence from Random Assignment of Students to Professors,” *Journal of Political Economy*, 2010, 118 (3), 409–432.
- Chandler, Michael Alison**, “Schools Pinched in Hiring,” *The Washington Post*, June 24 2007.
- Chetty, Raj, John N Friedman, and Jonah E Rockoff**, “Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates,” Working Paper 19423, National Bureau of Economic Research, Cambridge, MA 2013.
- , —, **and** —, “Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood,” Working Paper 19424, National Bureau of Economic Research, Cambridge, MA 2013.
- Chingos, Matthew M and Martin R West**, “Do More Effective Teachers Earn More Outside the Classroom?,” *Education Finance and Policy*, 2012, 7 (1), 8–43.

- Clark, Melissa A., Hanley S. Chiang, Tim Silva, Sheena McConnell, Kathy Sonnenfeld, Anastasia Erbe, and Michael Puma**, “The Effectiveness of Secondary Math Teachers from Teach For America and the Teaching Fellows Programs,” Technical Report, National Center for Education Evaluation and Regional Assistance, Institute of Education Sciences, U.S. Department of Education, Washington DC September 2013.
- Clotfelter, Charles, Elizabeth Glennie, Helen Ladd, and Jacob Vigdor**, “Would Higher Salaries Keep Teachers in High-Poverty Schools? Evidence From a Policy Intervention in North Carolina,” *Journal of Public Economics*, 2008, 92 (5-6), 1352 – 1370.
- Clotfelter, Charles T., Helen F. Ladd, and Jacob L. Vigdor**, “How and Why Do Teacher Credentials Matter for Student Achievement?,” Working Paper 12828, National Bureau of Economic Research, January 2007.
- Corcoran, Sean P., William N. Evans, and Robert M. Schwab**, “Changing Labor-Market Opportunities for Women and the Quality of Teachers, 1957-2000,” *The American Economic Review*, 2004, 94 (2), pp. 230–235.
- Dee, Thomas and James Wyckoff**, “Incentives, Selection, and Teacher Performance: Evidence from IMPACT,” Working Paper 19529, National Bureau of Economic Research October 2013.
- Dolton, Peter and Wilbert van der Klaauw**, “The Turnover of Teachers: A Competing Risks Explanation,” *Review of Economics and Statistics*, 1999, 81 (3), 543–550.
- Duflo, Esther, Rema Hanna, and Stephen P. Ryan**, “Incentives Work: Getting Teachers to Come to School,” *American Economic Review*, June 2012, 102 (4), 1241–1278.
- Figlio, David**, “Can Public Schools Buy Better-Qualified Teachers?,” *Industrial and Labor Relations Review*, 2002, 55 (4), pp. 686–699.
- Flyer, Fredrick and Sherwin Rosen**, “The New Economics of Teachers and Education,” *Journal of Labor Economics*, 1997, 15 (1), pp. S104–S139.
- Fryer, Roland G.**, “Teacher Incentives and Student Achievement: Evidence from New York City Public Schools,” *Journal of Labor Economics*, 2013, 31 (2), 373–407.
- , **Steven D. Levitt, John List, and Sally Sadoff**, “Enhancing the Efficacy of Teacher Incentives through Loss Aversion: A Field Experiment,” Working Paper 18237, National Bureau of Economic Research, July 2012.
- Gates, Bill**, “How Teacher Development Could Revolutionize Our Schools,” *The Washington Post*, February 28 2011.
- Glazerman, Steven, Hanley Chiang, Alison Wellington, Jill Constantine, and Dan Player**, “Impacts of Performance Pay Under the Teacher Incentive Fund: Study Design Report,” Report, Mathematica Policy Research, Washington DC 2011 2011.

- Glewwe, Paul, Nauman Ilias, and Michael Kremer**, “Teacher Incentives,” *American Economic Journal: Applied Economics*, July 2010, 2 (3), 205–227.
- Goldhaber, Dan, Betheny Gross, and Daniel Player**, “Teacher Career Paths, Teacher Quality, and Persistence in the Classroom: Are Public Schools Keeping Their Best?,” *Journal of Policy Analysis and Management*, 2011, 30 (1), 57–87.
- Goodman, Sarena F and Lesley J Turner**, “The Design of Teacher Incentive Pay and Educational Outcomes: Evidence from the New York City Bonus Program,” *Journal of Labor Economics*, 2013, 31 (2), 409–420.
- Gordon, Robert, Thomas J. Kane, and Douglas O. Staiger**, “Identifying Effective Teachers Using Performance on the Job,” Discussion Paper 2006-01, The Hamilton Project, April 2006.
- Hanushek, Eric A.**, “The Evidence on Class Size,” in Susan E. Mayer and Paul E. Peterson, eds., *Earning and Learning: How Schools Matter*, Washington, DC: Brookings Institution Press, 1999, pp. 131–168.
- **and Steven G. Rivkin**, “Teacher Quality,” in Eric Hanushek and Finis Welch, eds., *Handbook of the Economics of Education*, Vol. 2, Elsevier, 2006, pp. 1051–1078.
- Hess, Rick**, “Missing the POINT: Tomorrow’s Big Merit Pay Study Will Tell Us...Nothing,” *Education Week: Rick Hess Straight Up [blog]*, September 2010.
- Hiatt, Fred**, “How Bill Gates Would Repair Our Schools,” *The Washington Post*, March 30 2009.
- Holmstrom, Bengt and Paul Milgrom**, “Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design,” *Journal of Law, Economics and Organization*, January 1991, 7 (special issue), 24–52.
- Hoxby, Caroline M.**, “The Effects of Class Size on Student Achievement: New Evidence from Population Variation,” *The Quarterly Journal of Economics*, 2000, 115 (4), 1239–1285.
- Kane, Thomas J., Daniel F. McCaffrey, Trey Miller, and Douglas O. Staiger**, “Have We Identified Effective Teachers? Validating Measures of Effective Teaching Using Random Assignment,” Research Paper, Bill & Melinda Gates Foundation, Seattle, Washington January 2013.
- Kane, Thomas J, Jonah E Rockoff, and Douglas O Staiger**, “What Does Certification Tell Us About Teacher Effectiveness? Evidence from New York City,” *Economics of Education Review*, 2008, 27 (6), 615–631.
- Keane, Michael P. and Kenneth I. Wolpin**, “The Career Decisions of Young Men,” *Journal of Political Economy*, June 1997, 105 (3), 473–522.

- Klein, Joel, Michelle Rhee, Peter C. Gorman, Ron Huberman, Carol R. Johnson, Andres A. Alonso, Tom Boasberg, Arlene C. Ackerman, William R. Hite Jr., Jean-Claude Brizard, Jose M. Torres, J. Wm. Covington, Terry B. Grier, Paul Vallas, Eugene White, and LaVonne Sheffield**, “How to Fix Our Schools: A Manifesto by Joel Klein, Michelle Rhee and Other Education Leaders,” *The Washington Post*, October 10 2010.
- Krieg, John M.**, “Teacher Quality and Attrition,” *Economics of Education Review*, 2006, 25 (1), 13 – 27.
- Krueger, Alan B.**, “Experimental Estimates of Education Production Functions,” *Quarterly Journal of Economics*, 1999, 114 (2), 497–532.
- Lavy, Victor**, “Evaluating the Effect of Teachers’ Group Performance Incentives on Pupil Achievement,” *Journal of Political Economy*, December 2002, 110 (6), 1286–1317.
- Lazear, Edward P.**, “Performance Pay and Productivity,” *American Economic Review*, 2000, 90 (5), 1346–1361.
- , “Teacher Incentives,” *Swedish Economic Policy Review*, 2003, 10 (2), 179–214.
- Manning, Alan**, *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton, NJ: Princeton University Press, 2005.
- Muralidharan, Karthik and Venkatesh Sundararaman**, “Teacher Performance Pay: Experimental Evidence from India,” *Journal of Political Economy*, February 2011, 119 (1), 39–77.
- Murnane, Richard J. and David K. Cohen**, “Merit Pay and the Evaluation Problem: Why Most Merit Pay Plans Fail and a Few Survive,” *Harvard Educational Review*, 1986, 56 (1), 1–17.
- and **Randall J. Olsen**, “The Effect of Salaries and Opportunity Costs on Duration in Teaching: Evidence from Michigan,” *The Review of Economics and Statistics*, May 1989, 71 (2), 347–352.
- Neal, Derek and Diane Whitmore Schanzenbach**, “Left Behind By Design: Proficiency Counts and Test-Based Accountability,” *The Review of Economics and Statistics*, 2010, 92 (2), 263–283.
- Prendergast, Canice**, “The Provision of Incentives in Firms,” *Journal of Economic Literature*, 1999, 37 (1), 7–63.
- Ransom, Michael R. and David P. Sims**, “Estimating the Firm’s Labor Supply Curve in a ‘New Monopsony’ Framework: Schoolteachers in Missouri,” *Journal of Labor Economics*, April 2010, 28 (2), 331–355.

- Rivkin, Steven G., Eric A. Hanushek, and John F. Kain**, “Teachers, Schools, and Academic Achievement,” *Econometrica*, 2005, 73 (2), 417–458.
- Rockoff, Jonah E., Brian A. Jacob, Thomas J. Kane, and Douglas O. Staiger**, “Can You Recognize an Effective Teacher When You Recruit One?,” *Education Finance and Policy*, 2011, 6 (1), 43–74.
- Rothstein, Jesse**, “Teacher Quality in Educational Production: Tracking, Decay, and Student Achievement,” *Quarterly Journal of Economics*, 2010, 125 (1), 175–214.
- , “Review of ‘Learning About Teaching: Initial Findings from the Measures Of Effective Teaching Project’,” National Education Policy Center, Boulder, CO, 2011.
- , “Teacher Quality Policy When Supply Matters,” Working Paper 18419, National Bureau of Economic Research, September 2012.
- **and William J. Mathis**, “Review of Two Culminating Reports from the MET Project,” National Education Policy Center, Boulder, CO, January 2013.
- Sass, Tim R.**, “The Stability of Value-Added Measures of Teacher Quality and Implications for Teacher Compensation Policy,” Policy Brief 4, National Center for Analysis of Longitudinal Data in Education Research, November 2008.
- Springer, Matthew G., Dale Ballou, Laura Hamilton, Vi-Nhuan Le, J.R. Lockwood, Daniel F. McCaffrey, Matthew Pepper, and Brian M. Stecher**, “Teacher Pay for Performance: Experimental Evidence from the Project On Incentives in Teaching,” National Center for Performance Incentives at Vanderbilt University, 2010.
- Staiger, Douglas O. and Johah E. Rockoff**, “Searching for Effective Teachers with Imperfect Information,” *Journal of Economic Perspectives*, 2010, 24 (3), 97–118.
- Taylor, Eric S. and John H. Tyler**, “The Effect of Evaluation on Teacher Performance,” *American Economic Review*, December 2012, 102 (7), 3628–3651.
- Tincani, Michela**, “Teacher Labor Markets, School Vouchers and Student Cognitive Achievement: Evidence from Chile,” December 2011. Unpublished manuscript.
- von Wachter, Till, Jae Song, and Joyce Manchester**, “Long-Term Earnings Losses Due to Mass Layoffs During the 1982 Recession: An Analysis Using U.S. Administrative Data from 1974 to 2004,” April 2009. Unpublished manuscript.
- West, Martin and Matthew Chingos**, “Teacher Effectiveness, Mobility, and Attrition in Florida,” in Matthew G. Springer, ed., *Performance Incentives: Their Growing Impact on American K–12 Education*, Brookings Institution Press Washington DC 2009, pp. 251–272.


Wilson, Steven F., “Success at Scale in Charter Schooling,” *Education Outlook* 3, American Enterprise Institute March 19 2009.

Winters, Marcus A and Joshua M Cowen, “Who Would Stay, Who Would Be Dismissed? An Empirical Consideration of Value-Added Teacher Retention Policies,” *Educational Researcher*, 2013, 42 (6), 330–337.

— **and** —, “Would a Value-Added System of Retention Improve the Distribution of Teacher Quality? A Simulation of Alternative Policies,” *Journal of Policy Analysis and Management*, Summer 2013, 32 (3), 634–654.


Wiswall, Matthew, “Licensing and Occupational Sorting in the Market for Teachers,” 2007. Unpublished manuscript.

Figure 1: Objective probabilities of bonus receipt and tenure denial and average subjective expectations, by true ability percentile & experience




Notes: Left panel shows mean objective (solid line) and subjective (dotted and dashed lines) probability of receiving a bonus in any year two or more years in the future. Right panel shows objective and subjective probabilities of being denied tenure after year two. All probabilities are averaged over all teachers with true ability τ , ignoring any labor supply responses to the alternative contracts. Ability is rescaled to a percentile score (in the baseline distribution), on the x-axis.

Figure 2: Average equivalent variation of bonus and tenure contracts, by ability and experience




Notes: Series show the changes in base wage (w^0) under the baseline contract that would be needed to match the average value obtained by teachers at different points in the ability-experience ($\tau - t$) distribution under the performance-based contracts with wages set to fix the number of teachers employed. Ability is rescaled to a percentile score (in the baseline distribution), on the x-axis. For example, the point (95, +3.8) on the dashed line in the left panel indicates that under the bonus contract the average teacher with τ at the 95th percentile of the baseline distribution and one year of experience obtains a value equivalent to what would be obtained with a 3.8% salary increase under the single salary contract.

Figure 3: Effect of alternative contracts on number of new entrants to teaching, by ability percentile




Notes: Figure shows percentage change, relative to the single salary contract, in the number of new hires at each ability (τ) level under each of the performance-based contracts, when wages are set to fix the number of teachers employed. Ability is rescaled to a percentile score (in the baseline distribution), on the x-axis.

Figure 4: Effect of alternative contracts on average teaching career length, by ability percentile



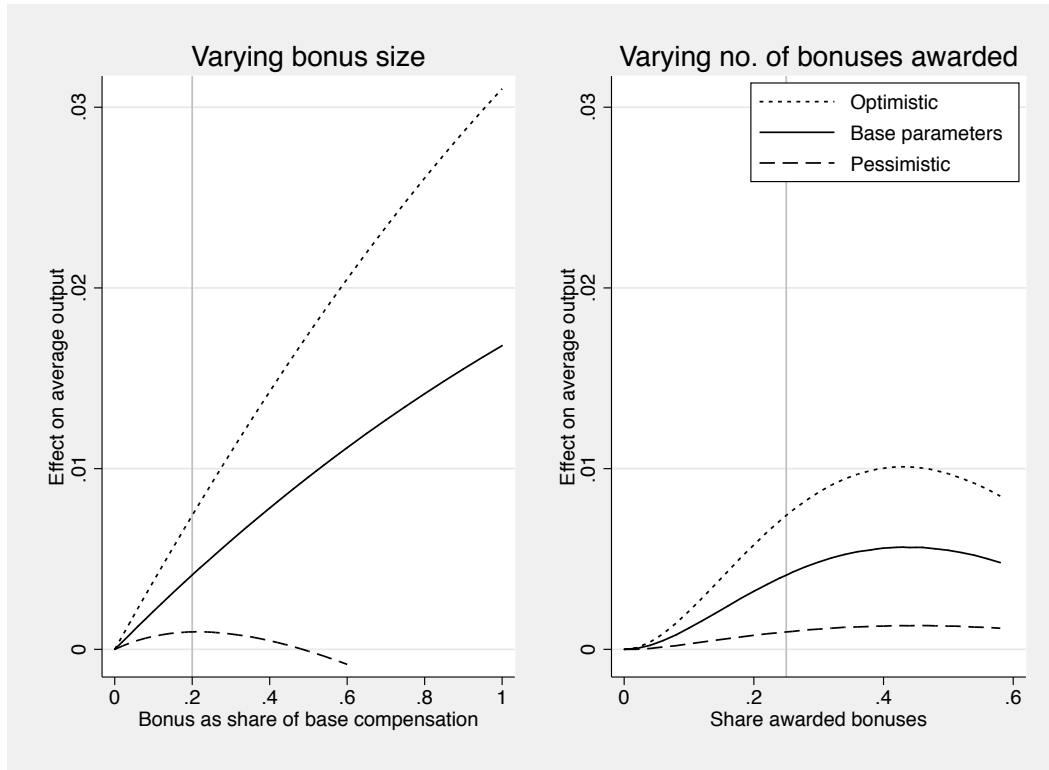
Notes: Figure shows percentage change, relative to the single salary contract, in the average career length at each ability (τ) level under each of the performance-based contracts, when wages are set to fix the number of teachers employed. Ability is rescaled to a percentile score (in the baseline distribution), on the x-axis.

Figure 5: Effect of alternative contracts on total number of teachers, by ability percentile




Notes: Figure shows percentage change, relative to the single salary contract, in the total number of teachers employed at each ability (τ) level under each of the performance-based contracts, when wages are set to fix the overall workforce size. Ability is rescaled to a percentile score (in the baseline distribution), on the x-axis.

Figure 6: Varying the parameters of the bonus contract



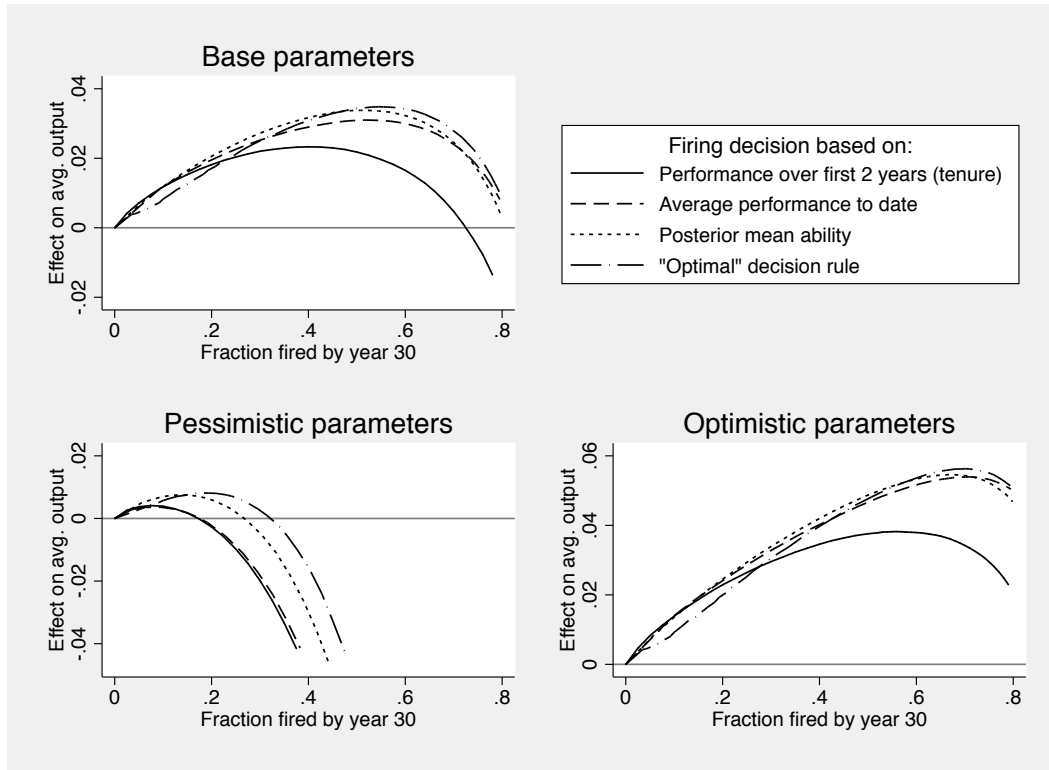
Notes: Panels show the change in average output, relative to the single salary contract and scaled in student-level standard deviations, associated with a shift to the bonus contract with the indicated bonus size (left panel) or threshold for bonus receipt (right panel, with x-axis scaled in terms of the share of current teachers who would receive bonuses each year). Parameters are as indicated in Table 1; base wages are assumed set to fix the total district budget. In left panel, bonuses are awarded to 20% of teachers (under the baseline ability distribution); in right panel, the bonus is set to 25% of base compensation. Vertical lines indicate contract parameters used for Table 3, Row 1.

Figure 7: Varying the tenure rate and timing of tenure decisions



Notes: Panels show the change in average output, relative to the single salary contract and scaled in student-level standard deviations, associated with a shift to the tenure contract with the indicated tenure denial rate, using different decision dates and parameter values. Parameters are as indicated in Table 1; base wages are assumed set to fix the total district budget. Marked points indicate the contract parameters (20% denied tenure, with decisions after the second year) used for Table 3, Row 1.

Figure 8: Comparing up-or-out tenure to ongoing retention decisions



Notes: Panels show the change in average output, relative to the single salary contract and scaled in student-level standard deviations, associated with shifts to contracts incorporating ongoing retention decisions. The decision rules are described in the text. X-axis measures the fraction of teachers who would be fired before the end of a 30-year career, given the ability distribution of current entering teachers. Parameters are as indicated in Table 1; base wages are assumed set to fix the total district budget.

Table 1: Values of key parameters

Parameter	Scenario		
	Baseline (1)	Pessimistic (2)	Optimistic (3)
<u>Effectiveness</u>			
σ_τ SD of teacher ability		0.15	
$r(t)$ Experience effect on productivity		-0.07 if $t=0$ -0.04 if $t=1$ -0.02 if $t=2$ 0 if $t>2$	
<u>Measurement</u>			
σ_ε SD of noise in annual performance		0.18	
<u>Teacher preferences & information</u>			
h Reliability of private information	0.1	0.05	0.25
δ Discount rate	0.97		
$u()$ von Neumann-Morganstern utility	Linear	CRRA, $\rho=3$	Linear
η Elasticity of entry w.r.t. w_0	1	0.5	1.5
ζ Absolute elasticity of exit w.r.t. w_0	1	0.5	1.5
λ_0 Annual exit hazard, base contract	0.08		
T Maximum career length	30		
κ Effect of being fired	1%* min($t,10\%$)	15%	1%* min($t,5\%$)
<u>Base contract</u>			
$g(t)$ Return to experience		0.015*t	
<u>Bonus contract</u>			
b Bonus size (as share of base pay)		20%	
s^B Share of teachers receiving bonus		25%	
α^B Base pay adjustment			
	Fixed quantity scenario	-3.5%	-2.7%
	Fixed budget scenario	-3.5%	-2.9%
			-3.5%
<u>Tenure contract</u>			
s^T Share of teachers receiving tenure	80%		
α^T Base pay adjustment			
	Fixed quantity scenario	+12.6%	+57.8%
	Fixed budget scenario	+10.4%	+27.1%
			+8.5%
			+7.4%

Table 2: Impact of bonus and firing contracts on teacher effectiveness and total costs

	Baseline	Inelastic demand		Fixed budget	
	(1)	Bonuses (2)	Tenure (3)	Bonuses (4)	Tenure (5)
<i>Outcomes under alternative contracts</i>					
Number of teachers	1	1.00	1.00	1.00	0.97
Teacher ability (τ)					
Mean	0.000	0.004	0.033	0.004	0.033
SD	[0.150]	[0.151]	[0.135]	[0.151]	[0.135]
Teacher experience					
Pct. 1st year	8.0%	8.0%	8.9%	8.0%	9.0%
Pct. 1st 3 years	30.9%	30.9%	31.1%	30.9%	31.6%
Mean	8.82	8.83	9.09	8.83	8.99
Teacher effectiveness ($\tau+r(t)$)					
Mean	-0.011	-0.007	0.021	-0.007	0.021
SD	[0.151]	[0.153]	[0.139]	[0.153]	[0.139]
Class size effect (rel. to baseline)	0	0.000	0.000	0.000	-0.014
Avg. output (teacher effect + class size)	-0.011	-0.007	0.021	-0.007	0.007
Salaries (expressed as multiple of baseline starting salary)					
Base starting salary	1.000	0.965	1.126	0.965	1.104
Average total pay	1.148	1.147	1.298	1.147	1.271
<i>Change from baseline</i>					
Number of teachers		n/a	n/a	+0.01%	-3.4%
Teacher ability (τ)					
Mean		+0.004	+0.033	+0.004	+0.033
SD		+0.001	-0.015	+0.001	-0.015
Teacher experience					
Pct. 1st year		-0.0 p.p.	+0.8 p.p.	-0.0 p.p.	+1.0 p.p.
Pct. 1st 3 years		-0.0 p.p.	+0.3 p.p.	-0.0 p.p.	+0.7 p.p.
Mean		+0.009	+0.270	+0.009	+0.167
Teacher effectiveness ($\tau+r(t)$)					
Mean		+0.004	+0.033	+0.004	+0.032
SD		+0.001	-0.012	+0.001	-0.012
Class size effect		+0.000	-0.000	+0.000	-0.014
Avg. output (teacher effect + class size)		+0.004	+0.033	+0.004	+0.018
Salaries					
Base starting salary		-3.5%	+12.6%	-3.5%	+10.4%
Average total pay		-0.1%	+15.0%	-0.0%	+12.3%

Notes: Simulations use baseline parameter values from Table 1, Column 1.

Table 3: Sensitivity of results to alternative parameters: Effects on teacher output (in student SDs)

Row	Bonuses			Tenure		
	Baseline	Pessimistic Alternative	Optimistic Alternative	Baseline	Pessimistic Alternative	Optimistic Alternative
	(1)	(2)	(3)	(4)	(5)	(6)
1 Baseline	+0.004	+0.001	+0.007	+0.018	-0.003	+0.023
Varying private information						
2 $h=0.05$	+0.004		+0.006	+0.018		+0.022
3 $h=0.25$	+0.005	+0.001		+0.019	-0.000	
Less noisy performance measure						
4 $\sigma_\epsilon=0.12$	+0.006	+0.001	+0.010	+0.023	+0.002	+0.028
Varying the supply elasticities						
5 $\eta=0.5$	+0.004		+0.006	+0.015		+0.019
6 $\eta=1.5$	+0.004	+0.001		+0.020	+0.004	
7 $\zeta=0.5$	+0.002		+0.004	+0.014		+0.019
8 $\zeta=1.5$	+0.006	+0.004		+0.021	+0.015	
9 ζ declines with t	+0.006	+0.004	+0.007	+0.020	+0.014	+0.023
10 η rises with μ	+0.004	+0.001	+0.007	+0.018	+0.002	+0.022
Other						
11 Risk averse (CRRRA=3)	+0.003		+0.005	+0.017		+0.022
12 High firing cost ($\kappa = -0.15$)				+0.015		+0.023
13 $\lambda=6\%$ in years 1-4; 3% thereafter	+0.003	+0.000	+0.005	+0.018	-0.005	+0.023

Notes: Row 1 uses parameter values from Table 1. Each successive row changes one parameter at a time, as indicated; entries where the alternative parameter matches that used in Row 1 are suppressed. All simulations assume a fixed budget, as in Columns 4-5 of Table 2. Table entries represent impacts of the alternative contracts on average teacher productivity, in student-level standard deviations relative to the baseline contract, incorporating changes in ability, in experience, and in class size. In Row 9, $\zeta = 1.5 * \min(0, t)$. In Row 10, $\eta = 1.01 + 0.0649(\mu/\sigma_\mu)$. This approximates the elasticity that arises in a Roy model with $corr(\mu, \omega_1) = 0.1$.

Table 4: Effects of contracts with fixed costs of entry

	No entry cost in baseline	Entry cost in baseline	
		Total impact	Incremental impact of performance contract
	(1)	(2)	(3)
Introduce bonus contract	+0.004	+0.004	
Introduce tenure contract	+0.018	+0.018	
Eliminate entry cost			
No other changes		+0.006	
and introduce bonus contract		+0.010	+0.004
and introduce tenure contract		+0.025	+0.018

Notes: All simulations assume a fixed district budget, as in Columns 4-5 of Table 2. Simulations in Column 2, Rows 1-2 assume that entering teachers must pay a fixed entry cost equal to w^0 , both under the baseline contract and under the alternatives. Table entries represent impacts of the alternative contracts on average teacher productivity, in student-level standard deviations relative to the baseline contract (with a fixed entry cost in Column 2, but without one in Column 1), incorporating changes in ability, in experience, and in class size.

Table 5: Optimal bonus and firing policies

	Baseline (1)	Pessimistic (2)	Optimistic (3)
<i>A. Bonuses based on two-year moving average performance</i>			
Varying bonus size, with baseline threshold			
Bonus size	100%*	21%	100%*
Impact	+0.0168	+0.0010	+0.0310
Varying threshold, with baseline size			
Fraction awarded bonus	43%	46%	43%
Impact	+0.0057	+0.0013	+0.0101
<i>B. Tenure policies</i>			
Varying date of tenure and fraction denied			
Year of tenure decision	3	5*	3
Fraction denied tenure	41%	13%	57%
Impact	+0.0241	+0.0062	+0.0383
Varying fraction denied, fixing decision in year 2			
Fraction denied tenure	40%	7%	56%
Impact	+0.0233	+0.0041	+0.0382
<i>C. Ongoing retention decisions</i>			
Based on average performance to date			
Fraction not retained	52%	10%	71%
Impact	+0.0310	+0.0036	+0.0540
Based on posterior mean			
Fraction not retained	52%	14%	67%
Impact	+0.0338	+0.0075	+0.0546
Based on optimal (static) decision rule			
Fraction not retained	55%	18%	70%
Impact	+0.0348	+0.0082	+0.0563

Notes: Asterisks represent boundary maxima – I do not consider bonuses larger than 100% of base pay or tenure decisions occurring after year 5. All simulations assume a fixed district budget. Impacts represent the net change in average teacher productivity, in student-level standard deviations relative to the baseline contract, incorporating changes in ability, in experience effects, and in class size.

A Appendices - Not for Publication

A.1 Search model

Each teacher draws a single outside job offer each year. If she accepts the offer, she exits teaching forever. The outside offer arrives after the teacher learns her previous year's performance (and is paid on that basis).

Outside offers are indexed by the continuation value that they provide, ω . I assume that the outside offer received prior to year $t > 1$, ω_t , has a censored Pareto distribution:

$$F_t(\omega_t) = \begin{cases} 0 & \text{if } \omega_t \leq V_t^0 \lambda_0^{1/\zeta'_t} \\ 1 - \lambda_0 \left(\frac{V_t^0}{\omega_t}\right)^{\zeta'_t} & \text{if } V_t^0 \lambda_0^{1/\zeta'_t} < \omega_t < HV_t^0 \\ 1 & \text{if } HV_t^0 \leq \omega_t. \end{cases} \quad (\text{A.1})$$

Here, V_t^0 is the continuation value if the teacher remains in teaching in year t under the baseline, single salary contract (which is constant across teachers), λ_0 is the annual exit hazard under this contract, and H is the maximum outside wage, expressed as a fraction of the inside continuation value.²⁶ Importantly, the distribution of ω_t is independent of the teacher's ability as a teacher, τ_i . Thus, as the teacher learns about τ_i she does not simultaneously learn about her future outside options (and vice versa).

Under the outside distribution (A.1), the probability that a teacher who would obtain continuation value $V_t \in [V_t^0 \lambda_0^{1/\zeta'_t}, HV_t^0]$ in teaching will instead exit is $\lambda_t(V_t) = \Pr\{\omega_t > V_t\} = \lambda_0 (V_t^0/V_t)^{\zeta'_t}$, with $\partial \ln \lambda_t(V_t)/\partial \ln V_t = -\zeta'_t$. The model in the main text is developed in terms of the negative of the elasticity of the exit hazard with respect to the inside wage under the baseline contract, $\zeta \equiv -\partial \ln \lambda_t/\partial \ln w^0 = -\partial \ln \lambda_t/\partial \ln V_t * \partial \ln V_t/\partial \ln w^0 = \zeta'_t * \partial \ln V_t/\partial \ln w^0$. The latter fraction varies with t . I thus solve recursively for this elasticity – which depends on $\zeta'_s, s > t$, but not on ζ'_t itself – and use it to define the elasticity parameter in (A.1) as $\zeta'_t \equiv \zeta * (\partial \ln V_t^0/\partial \ln w^0)^{-1}$.

The distribution of the initial non-teaching offer, ω_1 , is similar to that of offers later in the career, though here the shape parameter is computed as $\zeta'_1 \equiv \eta (\partial \ln V_1^0/\partial \ln w^0)^{-1}$.

²⁶The use of a censored distribution ensures that V_t is finite for any ζ'_t . It has no effect on the results so long as the censoring point is high enough that offers at that point are always accepted. I set $H = 2$, satisfying this criterion for the contracts under consideration here.

A.2 Solving the model

Equation (3) does not have a closed-form solution, but for any specified contract it can be solved recursively. Under the learning model developed above, the distribution of period- t performance measure given θ_{t-1} is

$$y_t | \theta_{t-1} \sim \mathbb{N} \left(\hat{\tau}_{t-1}, \frac{1}{\frac{1}{(1-h)\sigma_\tau^2} + \frac{t-1}{\sigma_\varepsilon^2}} + \sigma_\varepsilon^2 \right). \quad (\text{A.2})$$

This is a univariate distribution that can easily be computed for any specified value of $\hat{\tau}_{t-1}$. Given $\hat{\tau}_{t-1}$ and y_t , computation of $\hat{\tau}_t$ is trivial.

The recursive solution thus has three steps. First, I compute $w_T^C(y_1, \dots, y_T)$, the final period wage under contract C as a function of the performance signals to date. Second, I compute the value of remaining in teaching in period T , $V_T(\theta_{T-1}; C)$, as a function of θ_{T-1} , by integrating w_T^C over the conditional distribution of y_T given by (A.2). Third, for each $t < T$, given estimates of $V_{t+1}(\theta_t; C)$ as a function of θ_t , I compute $w_t^C(y_1, \dots, y_t)$ as a function of y_t and θ_{t-1} , then integrate each over the distribution of y_t (and therefore of θ_t) given θ_{t-1} to obtain $V_t(\theta_{t-1}; C)$.

The state space θ_t is of dimension $t + 1$, creating a dimensionality problem for careers of reasonable length. Note, however, that each of the contracts considered above reduces the state space for computation of w_t^C from the t -dimensional distribution $\{y_1, \dots, y_t\}$ to a one- or two-dimensional distribution: $\{y_{t-1}, y_t\}$ for the bonus contract and $\{\bar{y}_t\}$ for the tenure and alternative firing contracts. Meanwhile, the teacher's assessment of her own ability at the end of period $t - 1$ can be summarized either by the single variable $\hat{\tau}_{t-1}$ or by the pair $\{\mu, \bar{y}_{t-1}\}$. I can thus focus on state spaces of only two dimensions, $\tilde{\theta}_{t-1} = \{\hat{\tau}_{t-1}, y_{t-1}\}$ for the bonus contract or $\check{\theta}_{t-1} = \{\mu, \bar{y}_{t-1}\}$ for the tenure and firing contracts. I approximate the joint distributions of these two-dimensional state variables and y_t with grids of 149^3 points spaced to have equal probability mass.

Having computed $V_t(\theta_{t-1}, C)$ for each t , θ_{t-1} , and C , I simulate the impact of policies by drawing potential teachers from the $\{\mu, \tau\}$ distribution, then drawing performance measures $\{y_1, \dots, y_T\}$ for each. For each career, I compute θ_{t-1} and V_t at each year t , and use these to compute the effects of contract C on the probability of entering the profession and, conditional on entering, on surviving to year t . Note that I need not model the distribution of $\{\mu, \tau\}$ in the *population* of potential teachers – under my constant elasticity assumptions, changes in the returns to teaching induce proportional changes in the amount of labor supplied to teaching by each type that do not depend on the number of people of that type in the population.

A.3 Market clearing

Alternative contracts may yield greater or lesser entry or persistence in aggregate. For example, adding performance bonuses without reducing base pay will yield more entry from high- μ teachers and greater persistence of high- $\hat{\tau}_t$ teachers, without offsetting reductions from teachers with low μ or $\hat{\tau}_t$. Under each alternative contract, I compute the steady-state size of the teacher workforce, assuming that the contract has been in place for at least T years and that the same number of entering teachers have been hired in each year.

I consider two scenarios for labor demand. First, I assume that the education system will require the same number of teachers under the alternative contracts as are required under the baseline contract; where my computation yields a larger or smaller workforce than in baseline, I assume that the base salary is adjusted upward or downward to yield the appropriate number of teachers. The α^B and α^F parameters listed for the “fixed quantity” scenario in Table 1 are the adjustments required given the other parameters listed there; these are found via a numerical search algorithm. My second scenario assumes instead that the system’s total budget is fixed, so that increases in average teacher salaries must be offset by reductions in the number of teachers (and therefore by increases in class size). The “fixed budget” scenario rows in Table 1 show the α parameters that balance the district’s budget, given suitable changes in class size.

A.4 Optimal firing thresholds

In Section 5.2, I consider the optimal choice of thresholds (i.e., cutoff values of \bar{y}_t) for firing teachers at each year t . I compute these thresholds as the solution to the district’s dynamic optimization problem, assuming that the district pays a cost of firing a teacher that is proportional to the expected number of years remaining in the teacher’s career and that the district is myopic about potential labor supply responses. Specifically, let x_t represent the number of years that a teacher with t years of experience can be expected to remain in teaching given an annual exit probability of λ_0 and certain retirement after year T . It can be shown that

$$x_t = \frac{1 - \lambda_0}{\lambda_0} \left(1 - (1 - \lambda_0)^{T-t} \right).$$

Let $W_t(\bar{y}_t; c^{fire})$ represent the value of retaining a teacher (i.e., offering her employment for the next year) after year $t < T$ if her average performance to date is \bar{y}_t and the firing cost is c^{fire} ; let $W_0(c^{fire})$ represent the value of hiring a new teacher from the baseline ability distribution; and let $Z_t(c^{fire}) \equiv W_0(c^{fire}) - c^{fire}x_t$ represent the value of firing a teacher after year t . Then the continuation value of retaining a teacher after year $t = T - 1$ is:

$$W_t(\bar{y}_t; c^{fire}) = \lambda_0 W_0(c^{fire}) + (1 - \lambda_0) \left(\phi_t \bar{y}_t + r(t+1) + \delta W_0(c^{fire}) \right), \quad (\text{A.3})$$

where $\phi_t = \frac{\sigma_\tau^2}{t^{-1}\sigma_\varepsilon^2 + \sigma_\tau}$ and thus $E[\tau|\bar{y}_t] = \phi_t \bar{y}_t$. (δ is the discount rate.) The continuation value of retaining a teacher after year $t < T - 1$ is:


$$W_t(\bar{y}_t; c^{fire}) = \lambda_0 W_0(c^{fire}) + (1 - \lambda_0) \left(\phi_t \bar{y}_t + r(t+1) + \delta E \left[\max \left\{ W_{t+1}(\bar{y}_{t+1}; c^{fire}), Z_{t+1}(c^{fire}) \right\} \middle| \bar{y}_t \right] \right), \quad (\text{A.4})$$

Thus, the value of hiring a new teacher must be

$$W_0(c^{fire}) = 0 + r(0) + \delta E \left[\max \left\{ W_1(y_1; c^{fire}), Z_1(c^{fire}) \right\} \right]. \quad (\text{A.5})$$


Given a choice of c^{fire} and a hypothesized value for W_0 , one can use (A.3), (A.4), and (A.5) recursively to solve for the implied value of W_0 . The fixed point for this is the value $W_0(c^{fire})$. Moreover, the firing thresholds at year t are the values of \bar{y}_t that equate $W_t(\bar{y}_t; c^{fire})$ with $Z_t(c^{fire})$, and these can be used to compute the share of entering teachers who will be fired at some point in their careers. The estimates in Figure 8 are obtained by choosing a range of values for c^{fire} ; using a numerical search algorithm to find the fixed point W_0 given c^{fire} ; computing the firing thresholds implied by these values and the resulting firing shares; and then solving the labor supply model given these thresholds for the market-clearing wages and average productivity levels.

Figure A.1: Empirical one-year attrition hazards from the 1999/00 Schools and Staffing Survey/Teacher Follow-Up Survey




Notes: Solid line shows fraction of teachers at each experience level in the 1999-2000 Schools and Staffing Survey who are not teaching at the time of the one-year Teacher Follow-Up Survey. Dashed line codes as non-exits (a) individuals caring for family members at the time of the follow-up who say they plan to return to teaching within a year and (b) individuals who continue to work for state & local governments in non-teaching jobs in elementary and secondary education (e.g., as principals). Vertical line corresponds to the assumed retirement date ($T = 30$) used in simulations here. Horizontal lines correspond to the assumed annual attrition hazards used in the main (solid line; $\lambda = 0.08$) simulations and in the alternative simulation in Table 3, Row 13 (dotted line; $\lambda = 0.06$ for experience < 5 and $\lambda = 0.03$ thereafter). Sample sizes average 122 teachers per year of experience below 30.

Figure A.2: Probability of ever being fired over a 30-year career under different decision rules, by true ability



Notes: Firing probabilities are for teachers who remain in teaching for 30 years or until fired, whichever comes first. See text for definition of decision rules.

Figure A.3: Cumulative firing probability by true ability decile and experience under different decision rules



Panels show probability that a teacher will be fired on or before year t , given ability in the indicated group, assuming that no teacher exits voluntarily.